P517270 PDF [Page: 3 of 40]

Copy No. 12

DEFENCE R&D DÉFENSE

On-Screen Alarms in Computer-Aided Detection Systems

Combining Signal Detection, Human Factors, and System Design

Ronald T. Kessel

Defence R&D Canada

Technical Memorandum

DREA TM 2001-184

November 2001

National Defence Défense nationale

Canada

P517270.PDF [Page: 4 of 40]

Copy No. _____

On-Screen Alarms in Computer-Aided Detection Systems

Combining Signal Detection, Human Factors, and System Design

Ronald T. Kessel

Defence Research Establishment Atlantic

Technical Memorandum

DREA 2001-184

November 2001

P517270 PDF [Page: 5 of 40]

104-06987/

Author

Ronald T. Kessel

Approved by

Dr. Ron Kuwahara

Head/ Electromagnetics Section

Approved for release by

Kirk Foster

DRP Chair

© Her Majesty the Queen as represented by the Minister of National Defence, 2001

© Sa majesté la reine, représentée par le ministre de la Défense nationale, 2001

P517270.PDF [Page: 6 of 40]

Abstract

It is usual for computer-aided detection (CAD) systems to report automated alarms by highlighting their position on the operator's primary viewing screen. Here we carry out a "cost-benefit" analysis of the practice, contrasting its operational utility and liabilities. Because on-screen alarms are so very central to CAD systems, a cost-benefit analysis provides a framework for a critique of CAD systems more generally, combining aspects of signal detection theory, human factors, and system design. The analysis highlights the shortcomings of existing CAD systems, and identifies the source of widespread dissatisfaction with them, namely, why they are seen more often as a burden than a help by operators. The primary cause, it is argued, is that CAD systems are designed without a clear specification of the assisting role that automatic target recognition (ATR) is to fulfil for the operator. Clarified purpose in turn requires well-defined fusion rules by which the automated alarms are combined with the operator's, and by which the assisting goals are ultimately achieved. Given clarified purpose and fusion rules, it can be shown that the optimal integration of the ATR as an operator aid is likely to require radically different algorithm settings than if—as ATR developers generally assume for their algorithms—ATR were to operate autonomously. This may be a significant change in perspective for system designers who have so far avoided questions of human factors and overall performance. Non-conventional modes of CAD system operation are considered, showing how they offer different opportunities for assisting the operator.

Résumé

On utilise habituellement des systèmes de détection assistée par ordinateur (DAO) pour indiquer des alarmes automatisées en mettant en évidence leur position à l'écran principal de l'opérateur. Le document analyse la rentabilité de cette méthode et en fait ressortir les avantages et les inconvénients du point de vue opérationnel. En raison de l'importance cruciale des alarmes à l'écran pour les systèmes de DAO, l'analyse de rentabilité établit un cadre permettant de critiquer les systèmes de DAO dans une perspective générale, qui combine des aspects liés aux principes de détection des signaux, aux facteurs humains et à la conception des systèmes. L'analyse souligne les lacunes des systèmes de DAO existants, indique la source de la large insatisfaction qu'ils suscitent et précise pourquoi les opérateurs les considèrent souvent comme un fardeau plutôt que comme un moyen d'aide. Selon le document, cette situation tient essentiellement au fait que les systèmes de CAO sont conçus sans précision claire quant au rôle de soutien que la reconnaissance automatique des cibles (ATR) devrait jouer pour l'opérateur. La clarification de ce rôle exige quant à elle des règles de fusion bien définies, par lesquelles les alarmes automatisées sont combinées à celles de l'opérateur et grâce auxquelles les objectifs de soutien sont finalement atteints. Ce rôle ayant été clarifié et des règles de fusion ayant été établies, on peut démontrer que l'intégration optimale de l'ATR comme moyen d'aide à l'opérateur devrait exiger des algorithmes

517270 PDF [Page: 7 of 40]

radicalement différents de ceux, généralement adoptés en développement d'ATR, qui seraient utilisés si le système ATR devait fonctionner de façon autonome. Il peut s'agir là d'un changement de perspective important pour les concepteurs de systèmes qui ont jusqu'à maintenant évité les questions de facteurs humains et de rendement global. Le document examine des modes non classiques de fonctionnement d'un système de DAO et montre comment ils offrent des possibilités différentes d'aide à l'opérateur.

P517270 PDF [Page: 8 of 40]

Executive summary

Background

Sidescan sonars are the sensor of choice for mine countermeasures (MCM) when large areas of the seafloor must be inspected for mines and mine-like clutter. DRDC projects in MCM are directed in part towards the automated detection of mines in sidescan imagery, as an aid to the operator who must visually inspect sonar imagery for mine-like objects. Thus, the Remote Minehunting (RMS) Technology Demonstrator (TD) project underway at DREA includes a significant computer-aided detection (CAD) component among its priorities. The present work focuses on early design questions in support of that project, especially how automatic target recognition (ATR) ought to be integrated as a productive operator aid, to ensure that the combined human-machine performance is better than that of the operator working alone.

Principal results

The emphasis in much CAD system development has so far been towards improved ATR algorithms. It is shown here that shortcomings in CAD are also due in part to the lack of clear objectives for the assisting role of the ATR, and of systematic rules (fusion rules) by which the ATR's alarms are to be handled—both of which determine the mode of human-machine interaction. The operator-assisting goals of ATR, and four possible fusion rules are reviewed here. It is shown how the conventional CAD system concept is likely to burden, rather than help the operator, unless its natural advantage (increased operator vigilance) is emphasized from the outset. And it is shown how alternative fusion rules offer benefits in other respects (reduced workload, and added information), and therefore warrant consideration.

Significance of results

The supporting analysis illustrates the performance trade-offs of alternative modes of CAD system operation, and demonstrates their potential to help or hinder the operator.

Future Work

As ATR algorithms improve with development over time, the analysis can be used to turn growing ATR autonomy to best advantage. Many of the concepts touched on here will be used to guide the development of CAD for the RMS TD project underway at DREA.

Ronald T. Kessel. 2002. On-Screen Alarms in Computer-Aided Detection Systems. DREA TM 2001-184. Defence Research Establishment Atlantic.

517270.PDF [Page: 9 of 40]

Sommaire

Contexte

Les sonars à balayage latéral sont des capteurs de choix pour la lutte contre les mines (LCM) lorsqu'il faut examiner de grandes zones du fond marin pour découvrir les mines et les fouillis d'échos causés par des clutter objets semblables aux mines. Les projets en LCM de RDDC sont partiellement consacrés à la détection automatisée des mines par imagerie à balayage latéral, comme aide à l'opérateur qui doit inspecter visuellement l'imagerie sonar pour détecter les objets semblables aux mines. Ainsi, le projet de démonstrateur de technologie (DT) du système télécommandé de chasse aux mines (STCM), actuellement en cours au CRDA, englobe au nombre de ses priorités une importante composante de détection assistée par ordinateur (DAO). Les travaux actuels portent sur les premières questions de conception à l'appui de ce projet, particulièrement sur la façon dont la reconnaissance automatique des cibles (ATR) devrait être intégrée comme moyen productif d'aide à l'opérateur, et ils visent à maintenir un rendement combiné humain-machine meilleur que si l'opérateur travaillait seul.

Principaux résultats

Une large part des travaux de développement des systèmes de DAO ont jusqu'à maintenant cherché à améliorer les algorithmes d'ATR. Ce document montre que les lacunes de la DAO sont aussi partiellement dues au manque d'objectifs clairs quant au rôle d'appui de l'ATR ainsi qu'à l'absence de règles systématiques (règles de fusion) par lesquelles les alarmes de l'ATR doivent être traitées, ces deux facteurs déterminant le mode d'interaction humain-machine. Les objectifs d'aide à l'opérateur de l'ATR ainsi que quatre règles de fusion possibles sont examinés. Il est démontré comment le principe du système ordinaire de DAO pourrait constituer un fardeau plutôt qu'un moyen d'aide pour l'opérateur, à moins que l'accent soit mis dès le départ sur son avantage naturel soit, une (vigilance accrue de l'opérateur). On indique aussi comment des règles de fusion différentes offrent des avantages à d'autres égards (réduction de la charge de travail et augmentation de l'information) et pourquoi elles méritent donc d'être examinées.

Importance des résultats

L'analyse illustre les compromis en matière de rendement qui sont liés aux divers modes de fonctionnement d'un système de DAO, et elle évalue la possibilité qu'ils aident l'opérateur ou qu'ils constituent un fardeau.

Plans d'avenir

À mesure que les algorithmes d'ATR s'amélioreront avec le développement, l'analyse pourra servir à tirer le meilleur parti de l'autonomie croissante des systèmes d'ATR. Un

P517270 PDF [Page: 10 of 40]

bon nombre des principes abordés dans le document serviront à orienter le développement de la DAO dans le cadre du projet de DT du STCM, actuellement en cours au CRDA

Ronald T Kessel 2002 French Title Here DREA TM 2001-184 Centre pour la Recherche de la Défence Atlantique

P517270.PDF [Page: 11 of 40]

Table of contents

Abstract i

Résumé i

Executive summary iii

Sommaire iv

Table of contents vi

List of figures viii

Acknowledgements x

1. Introduction 1

2. Goals, costs, and benefits 2

2.1 Overall system goals 2

2.2 Goals of on-screen alarms 3

2.2.1 Alert and inform 3

2.2.2 ATR feedback 3

2.2.3 Operator vigilance 4

2.3 The costs of on-screen alarms 4

2.3.1 ATR adjustments 4

2.3.2 Alarm masking 5

2.3.3 Operator bias 5

2.3.4 Undermined operator confidence 5

3. ATR mode of operation 6

4. The signal detection theory (SDT) 7

5. Combined human and computer performance 8

5 1 Joint performance 11

5.2 Operator Workload 12

P517270.PDF [Page: 12 of 40]

5.3 Operator Vigilance 14
6. Example 15
6.1 Bivariate Gaussian distribution 15
6.2 CAD system performance 16
6.3 Operator Workload 18
7. Conclusions 18
References 21

P517270.PDF [Page: 13 of 40]

List of figures

1 The statistical model of a detector entails two conditional probability density functions: one for targets (1), and another for non-targets (0). An alarm is registered whenever the decision variable x exceeds the detection threshold x_T. The value of the threshold depends on larger search-mission factors, such as the prior probability of encountering a target, and the costs and risks associated with detection errors (missed targets, or false alarms). 8

2 The full range of a detector's performance is given by its receiver-operator characteristics. The detection threshold determines the point at which a detector operates on this curve; high thresholds tending towards the lower left, and low thresholds to the upper right of the graph. The quality of the detector depends on its sensitivity (class separation in Fig.(1)), with higher sensitivities giving higher probabilities of detection for a given probability of false alarm. 9

3 The combined ATR (A) and human (H) detection can be modeled using bivariate random distributions; one for the target class (1), and another for non-targets (0). Here the ATR decision variable is plotted on the horizontal axis, and the human operator's (effective) decision variable is plotted on the vertical. Each agent registers an alarm when their decision variable exceeds their threshold. The two thresholds divide the plane into four quadrants, labeled A, B, C, and D here. . . . 10

4 ROC curves for equal operator and ATR detection assuming the bivariate normal distributions given in (11) and (12). The plain solid line is the ROC curve for the operator and ATR operating independently, assumed to be equal here for simplicity. The squares mark the best possible performance under fusion rule 3, the "OR" rule. The circles mark the best possible performance under rule 4, the "AND" rule. In this case rule 4 clearly outperforms rule 3. But this is not generally true. The reverse may be true, depending on the disposition of the class distributions (Fig.(3)). Note that the optimal performance under either rule is only achieved when the operator and the ATR adjust their thresholds in concert. The optimization was achieved here by numerical iteration. It is not clear how the optimal would be achieved in practice, so these curves represent the ultimate performance limits under ideal conditions, for the assumed density functions. . . 17

5 The relative change in operator workload under fusion rule 2, as a function of the ATR detection threshold. Negative changes signify reductions in workload. In this case, the anticipated change in operator workload would not be very significant, less than 5 % for realistic ATR thresholds greater than 1.0. 19

6 The relative change in operator workload under fusion rule 3, as a function of the ATR detection threshold. Unlike rule 2, this rule always reduces operator workload. In this example, the reduction is more significant as well, perhaps a 10 % reduction for ATR thresholds greater than 1.0. 19

P517270.PDF [Page: 14 of 40]

7 The relative change in operator workload under fusion rule 4, as a function of the ATR detection threshold. This rule affords the greatest potential for workload reduction, perhaps a 75 % reduction for ATR thresholds near 1.0. A more detailed analysis would be required to determine if the performance at this workload reduction would be acceptable. Even a 20 % reduction near ATR thresholds of -1.0 would be greater than those encountered with the other fusion rules in this example. 20

P517270.PDF [Page: 15 of 40]

Acknowledgements

Discussions about human-machine performance with Dr. Nancy Allen at the Defence Research Establishment Atlantic have been instrumental in the development of this critique.

1. Introduction

In remote sensing for military applications, the problem of acquiring new and, in many ways, unpredictable targets, has so far proved too difficult to entrust to a computer alone. The question then arises whether computer aided detection (CAD) systems might be developed to assist, rather than to replace, the human operator. Such developments are in fact underway for sea minehunting. The side-scan sonar echographs recorded in sea minehunting look rather like aerial photographs of the seafloor. Both the computer and the human operator scan the images for mine-like objects, and the computer alerts the operator to its alarms by brightly marking their position on the operator's primary viewing screen. Much the same is being done elsewhere for air and ground targets using radar and infrared.

The reason for automation in such cases is of course to improve the overall detection performance above that of the operator working alone, without automation. Otherwise one would do precisely that—forego the automation. Thus it is the CAD developer's responsibility to show how, exactly, the automation is going to improve matters. Such work has often been upstaged by the design of the automatic target recognition (ATR) algorithms themselves, which are the "brains" of the automation. But until those algorithms advance to autonomous operation, the mode of human-machine interaction, the "body" of the automation, will continue to be equally important for the operational success of a CAD system.

The change in emphasis for ATR, from autonomous operation to operator aid, is more complicated than it first appears. Nowhere is this more evident than in an analysis of on-screen alarming. Being the dominant mode of communication between the computer and the human operator, on-screen alarms serve as the fulcrum by which ATR is leveraged for better or worse. For just as a misplaced fulcrum can make a lever's action heavier rather than lighter, so a haphazard approach to on-screen alarming can make the overall performance worse rather than better, leaving the operator little choice but to switch the automation off altogether. That is the risk CAD designers face.

Our purpose, then, is to present a framework in which such system-level issues can be addressed in an objective way, by weighing the operational benefits of on-screen alarming against its costs or penalties. We begin by reviewing the purposes of the CAD system as a performance-enhancing aid, as well as the more particular objectives of on-screen alarming as a sub-system. Then we use the signal detection theory (SDT) borrowed from monitored systems in human factors [1][2][3], which suits CAD systems particularly well, to draw implications for the design, not only of on-screen alarming, but of the larger CAD system, including the ATR. The analysis is kept general and qualitative, without specifying the sensor or signal quality, the particular ATR algorithms, the operator's expertise, and so forth, but it could be extended for more detailed analysis of particular CAD systems if experiments were first carried out to quantify their determining parameters.

Our scope is limited to CAD systems of the type used in sea minehunting and

P517270.PDF [Page: 17 of 40]

elsewhere, for which:

- automated alarms are marked directly on the operator's primary viewing screen, by somehow highlighting the target in question during the operator's first pass through the search;
- the signals are neither too complex nor too vast for routine visual inspection by a human operator in the absence of automation—i.e., operator's capacity is not overwhelmed from the outset, in the absence of ATR;
- the operator and the ATR algorithms monitor signals from the same sensor, though each possibly "seeing" a different dynamic range and resolution in the data;
- genuine targets are rare events amid highly variable clutter;
- the alarms raised by the human-machine combination may or may not be passed on to another higher level of target classification[1].

Other types of CAD systems, such as in which the quantity of signals prevents the operator from attending to more than a fraction of them, might also be treated using a modified form of the present analysis. ATR and on-screen alarms are just one of many performance-enhancing computer aids. Data management and mapping tools for remote sensing, providing convenient access to multi-resolution, multi-look, or multi-aspect data, are of course very important, but are not considered here.

A discussion of human factors finally requires that certain assumptions be made about the human operator. In taking up questions of ultimate CAD system performance, we assume that the human operator is expert and fully trained, and does not, as a novice operator might, rely on the ATR for their detection training. The possibility of using a CAD system to teach novice operators about the kinds of target-like signatures to look for in their viewing screen is not considered here, but that training potential should not be ignored in a larger operational study.

2. Goals, costs, and benefits

Before a detailed analysis is possible, we must clarify and the goals of the overall CAD system, and those of the on-screen alarm sub-system.

2.1 Overall system goals

A CAD system can improve the overall detection performance in a combination of ways, by

[1] A second level of alarm analysis often occurs quite naturally in practice because the operators ordinarily do not work in isolation. Several experienced operators and officers may be on hand for consultation. The implications of second-level alarm screening are important. A CAD system can usually tolerate a higher false alarm rate, for instance. But the implications are more quantitative than qualitative for our purposes, and are not considered here.

a)providing an independent threat assessment, much as if a second, untiring operator were scanning the image;

b)reducing the workload of the human operator, or

c)improving the vigilance of the human operator.

A CAD system that failed in all three respects would certainly be ineffectual, if not detrimental. These goals must guide the design at every stage.

Goal (*a*) is achieved in two ways. The ATR provides "backup" coverage if the operator has lapses of attention owing to overload, distraction, or fatigue. And by drawing on different, more quantitative target features, the ATR furthermore provides a more or less independent threat assessment of value in its own right. To speak of ATR independence is of course the first step towards making it autonomous—a point to which we shall return several times. In goal (*b*), the operator's total effort is reduced, presumably because the ATR is doing part of their task, reliably scanning signals and registering alarms. In goal (*c*), "improved vigilance" means increasing the arousal and sensitivity of the operator insofar as an on-screen alarm serves as a momentary diversion for the operator, increasing their sensitivity and vigilance as shown in human factors [2][3]. This is of course a particular goal of on-screen alarms in its own right, to which we therefore return in section 2.2.3.

2.2 Goals of on-screen alarms

Being a sub-system, on-screen alarming may have goals that are not always directly related to those of the overall system, and may in fact be at odds with them. Here we briefly review the purposes of on-screen alarms in their own right.

2.2.1 Alert and inform

Their main purpose is of course to alert the operator to a potential threat, while at the same time immediately conveying the target position and signature. One can hardly imagine a more natural way of doing this than highlighting the alarm directly on the operator's view of the signal as in on-screen alarming. But even good ideas may come with a cost, as we shall see.

2.2.2 ATR feedback

ATR developers have often found that, in order to accommodate the wide range of operating conditions typically encountered between missions, or even during a single mission, the sensitivity of the ATR algorithms must often be adjusted by the operator to eliminate excessive false alarms in some cases, or unresponsiveness in others. In sea minehunting, for instance, adjustments would typically be required in response to changes in seafloor clutter, texture,

P517270.PDF [Page: 19 of 40]

brightness and so forth [4][5]. On-screen alarms therefore play an important feedback role for the operator, indicating when ATR is working optimally—neither raising too many false alarms, nor missing obvious target-like signatures—, or when it is in need of adjustment. Oddly enough, this second purpose of on-screen alarming directly serves the automation itself, but overall performance only indirectly. The workload associated with such adjustments is included in the discussion of costs below.

2.2.3 Operator vigilance

A third purpose of on-screen alarming reiterates system-goal (*c*) above: to improve operator vigilance. The occasional on-screen alarm serves as a momentary diversion for the operator, raising their arousal and sensitivity, and serving as a reminder of what target-like signatures look like, all of which combats the decrease of the operator's vigilance over time. The key here is that the alarms are sufficiently rare to strike the operator as exceptional. This would not be true if the ATR algorithms raised an endless stream of false alarms, owing to improper adjustment for instance. A long absence of ATR alarms, on the other hand, would allow vigilance to lapse more than it otherwise might, had the occasional alarm been raised. Thus there is an optimal alarm rate maximizing vigilance. It has even been suggested that simulated false-alarms might deliberately be injected into the operator's signal during extremely low event periods, in order maintain that rate [3].

2.3 The costs of on-screen alarms

The operational costs of on-screen alarms potentially outweigh the benefits offered by automated detection. And the balance of costs against benefits is likely to change from mission to mission, or from time to time during a single mission, with changing clutter levels, environmental conditions, perceived risk, and so forth. The operator will be inclined to switch the ATR alarms off whenever the operational costs outweigh the benefits.

The costs enumerated in this section are difficult to quantify, in absolute or relative terms. Many depend on system and environmental particulars which cannot be pursued here, except to be itemized as unspecified, possibly significant costs implicit to on-screen alarming. The workload associated with the operator's response to an alarm, on the other hand, for which a more analytic treatment is possible, is deferred until section (5.2).

2.3.1 ATR adjustments

These adjustments are the cost of managing the automation itself, mentioned above in section (2.2.2). The operator's continual questioning whether new adjustments are required, and then the attention and effort required to make

those adjustments, can be a significant supervenient workload, particularly when they call for specialized knowledge of the ATR algorithms themselves. ATR designers naturally minimize the need for repeated on-site adjustments throughout a mission, though with varying degrees of success. A review of twelve existing CAD systems (mostly proprietary) concluded that "it is reasonable to expect [that] the operator may tune the selected algorithm on the order of, or less than, once per minute, for the type of ocean bottom or ambient conditions encountered" [5, Section 4.3.2].

2.3.2 Alarm masking

Alarm *masking* is the obscuring of the operator's primary screen by the ATR alarms themselves inasmuch as the alarms over-write or distort portions of the image that the operator may wish inspect more closely—the alarm of one prospective target overlaying the signature of another. The CAD system should therefore allow the operator to instantly toggle ATR alarms on and off, as most systems do in fact. Whenever the alarms are enabled, alarm masking is a variable cost of operation. It can be reduced somewhat by joining overlapping alarms into a single alarm, highlighting the boundary of larger regions of interest where several nearby alarms are better grouped together under a single alarm. Signaling alarms by local screen color changes, rather than by boldly circling them with solid boundaries, may also reduce masking. But this still leads to visual distortion because the intensity rendering of the screen (pixel brightness), and the perception of intensity by the human eye, both depend in part on color [6], hence changing colour produces visual changes in the target's strength and contrast.

2.3.3 Operator bias

Less obvious are the costs incurred precisely because an on-screen alarm draws the attention of the operator toward the alarm itself, and away from other portions of the screen. Thus the operator may miss a genuine target that eluded the ATR, simply because the operator's attention was drawn elsewhere by ATR false alarms. In other words, on-screen alarms bias the operator away from non-alarmed portions of the screen. This could hardly be otherwise, yet the system designer should consider minimizing the bias, possibly using a two-step approach, in which ATR alarms are withheld until after the operator has inspected the image independently, or by training the operator against such bias.

2.3.4 Undermined operator confidence

If false alarms occur at a high rate, then it is inevitable that many of them will appear obviously false to the operator, causing them to adopt an instinctively dismissive stance towards its alarms. This loss of confidence in the automated

aid, or possibly the failure to establish confidence in the first place, is equivalent in many ways to a reduction of operator vigilance regarding on-screen alarms. The majority of on-screen alarms should therefore challenge the operator, giving them pause to consider the target-like nature of the signature. Only a small minority should be obviously false at a glance. This is a function of the quality of the ATR and its settings.

3. ATR mode of operation

A more analytic approach to on-screen alarms requires that we be absolutely clear about the ATR's mode of operation within the larger CAD system—i.e., its degree of independence, the internal trade-off it makes between its probability of detection and probability of false alarm, and the operator's instructions for dealing with its alarms. The mode of operation is determined mainly by the way in which the alarms raised by the two detection agents—operator and ATR—are to be fused into a single alarm (or into final list of alarms compiled over time during a remote survey). Four such fusion rules are considered here:

1. *transient*—ATR alarms are deemed false unless actively acknowledged as true by the operator;
2. *durable*—ATR alarms are deemed true unless actively dismissed by the operator;
3. *Boolean "OR"*—the fused alarms include the union of all alarms raised, irrespective of the agent who raised them;
4. *Boolean "AND"*—the fused alarms include just those alarms raised by both the ATR and the operator.

The first is the most common, if not the only rule followed in existing CAD systems, presumably owing to the cautious stance usually adopted towards ATR—its alarms being completely subordinate to the operator, and in effect false by default. The caution is no doubt justified so long as the ATR's performance remains unknown, or is known to have failings.

The second rule confers slightly more independence to the ATR, but only insofar as the operator becomes overloaded and cannot inspect every ATR alarm, in which case all unverified alarms are presumed genuine by default, each launching all that a genuine alarm from the CAD console entails for the mission. At first glance this default may seem intolerable. Given a highly selective ATR, however, for which false alarms are in fact very rare due to its high detection *threshold* or high *sensitivity* (terms defined below), this second rule may become feasible as an operator aid and warrants investigation.

The third and forth rules assume unprecedented independence for the ATR. They could in fact be implemented mechanically, with no on-screen alarms whatsoever for the

operator. As we shall see, these rules are naturally inclined towards extreme ATR settings, that is, very low false alarm rates for "OR" rule 3, and very high false alarm rates for "AND" rule 4—"extreme" here meaning relative to the ATR settings if it were required to operate autonomously.

Rules 2 through 4 are radical insofar as they give the ATR some degree of independence apart from the operator. It is only through such independence that independent-agent performance gains—goal (*a*) above—could be realized, as we shall see.

4. The signal detection theory (SDT)

The mechanisms underlying target recognition by a human operator are very complicated to model in detail. Considered statistically, however, the net result target/non-target decision process resembles that of conventional detection theory. At least this is the starting point of the signal detection theory of monitored alarms in human factors [1] [3]. In conventional detection theory [7] [8] it is assumed that the real world presents a physical sensor with two classes of objects: non-targets (class 0) and targets (class 1) plus noise. The sensor in turn delivers the resultant signals to the detector (here the ATR or human operator), which then raises alarms with a probability of detection P_1 (the conditional probability of registering an alarm given a target), and with a probability of false alarm P_0 (the conditional probability of registering a false alarm given no target). These performance probabilities could be measured experimentally, by measuring the relative frequency of missed and false detections in a dummy target field for instance.

The performance probabilities are in turn attributed to probability density functions (p.d.f.'s), $p_1(x)$ and $p_0(x)$, defined on a single decision variable x, which triggers an alarm whenever x exceeds some threshold x_T; i.e.,

$$P_{1,0} = \int_{x_T}^{\infty} p_{1,0}(x)\, dx, \tag{1}$$

as shown in Figure (1). The justification for density functions in a single real variable x, while ignoring the detector's inner complexity, is given for automated systems by theoretical proofs in pattern recognition [9], and for humans by experiment [1] [3]. For our purposes, it is enough to infer an equivalent decision-variable representation of the form (1) from the performance probabilities, whether or not we are able to locate such a decision variable physically, or directly measure it somewhere inside the detector. Both the human operator and the ATR can be modeled in this way, each with their particular performance probabilities, decision variables, and thresholds. The superscripts ^{H} and ^{A} will therefore be assigned to variables, as required, in order to distinguish between parameters of human (H) and ATR (A) performance respectively.

The main advantage of the SDT, then, is its simplification of the decision-making process in each agent, treating them as random processes whose properties can be

P517270.PDF [Page: 23 of 40]

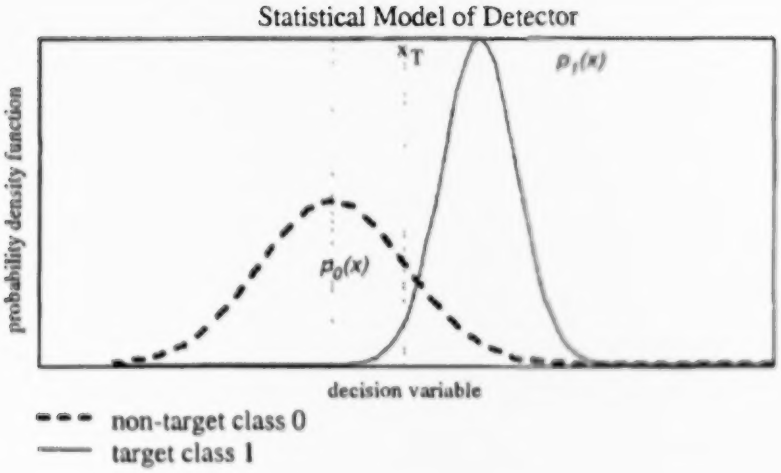

Figure 1: *The statistical model of a detector entails two conditional probability density functions: one for targets (1), and another for non-targets (0). An alarm is registered whenever the decision variable x exceeds the detection threshold x_T. The value of the threshold depends on larger search-mission factors, such as the prior probability of encountering a target, and the costs and risks associated with detection errors (missed targets, or false alarms).*

adjusted to suit the full range of detection problems, from "easy" to "difficult", or what amounts to the same thing, a full range of detector abilities, from low quality discriminators to high. The ability of a detector to discriminate between classes is called its *sensitivity*, and is determined by the separation between the two classes, targets and non-targets. High sensitivity, hence good performance, is achieved when there is little overlap between the distributions $p_1(x)$ and $p_0(x)$.

The balance struck between correct detection and false alarms depends furthermore on the *detection threshold* x_T, whose optimal value depends on mission factors beyond the signals themselves, such as the prior likelihood of encountering a target, and the operational costs and risks associated with making detection errors [7]. It is usual to consider a continuous range of possible thresholds when assessing a detector's performance, by plotting the receiver-operator characteristics (ROC) for the detector (see Fig.(2)), each point of which corresponds to a unique threshold x_T.

5. Combined human and computer performance

Figure (3) illustrates the joint probability density for the decision variables of the two agents. This is the two-dimensional extension of Figure (1), from one detection agent, to two. The human decision variable x^H runs along the vertical axis, the ATR x^A along

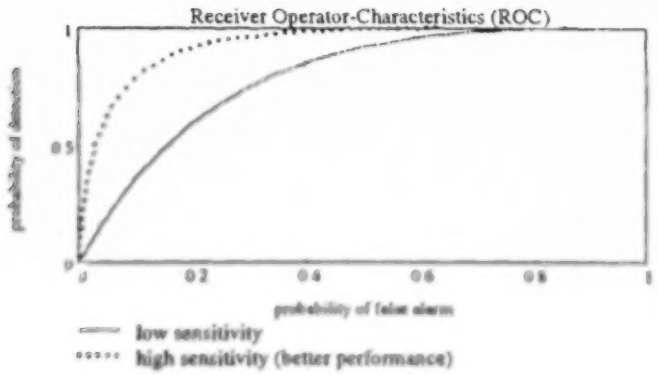

Figure 2: The full range of a detector's performance is given by its receiver-operator characteristics. The detection threshold determines the point at which a detector operates on this curve, high thresholds tending towards the lower left, and low thresholds to the upper right of the graph. The quality of the detector depends on its sensitivity (class separation in Fig (1)), with higher sensitivities giving higher probabilities of detection for a given probability of false alarm.

the horizontal axis, and their joint probability density now runs along a third axis out of the page (not shown) whose values are represented using contour lines. A single contour for each class is shown, encircling the regions where the two decision variables (x^A, x^H) tend to cluster when both agents view the signature of the same object. Here again, an easy detection problem, in which good target/non-target discrimination can be achieved without difficulty, is one in which the classes are well separated—i.e., the contours encircling large, equal percentiles of the each class do not overlap.

The decisions of the operator and ATR may be correlated. That is, the ATR may have a tendency to ascribe high x^A for the same objects to which the operator ascribes high x^H. Such correlation tilts the contours along a positive slope as shown in Figure (3). Strong positive correlation means that the ATR closely replicates the decisions of the human operator. Weak correlation indicates that the ATR is focusing on different features than the visual cues used by the human operator. The nature of the correlation has important implications for workload reduction and performance, as we shall see, mainly because it affects the disposition of the classes in Figure (3).

Pattern recognition theory considers how to draw decision boundaries in Figure (3) for optimal performance, maximizing P_1 for a given P_0 for instance [8] [9] [10]. CAD systems are unfortunately more constrained than pattern recognition problems generally are, in part because the operator and ATR ordinarily make binary target/non-target decisions, as opposed to ranking objects by assigning a number, or threat level to them, and in part because the operator's decision variable x^H is not available, both of which make the calculus of globally optimal decision boundaries impossible. It is possible for

P517270.PDF [Page: 25 of 40]

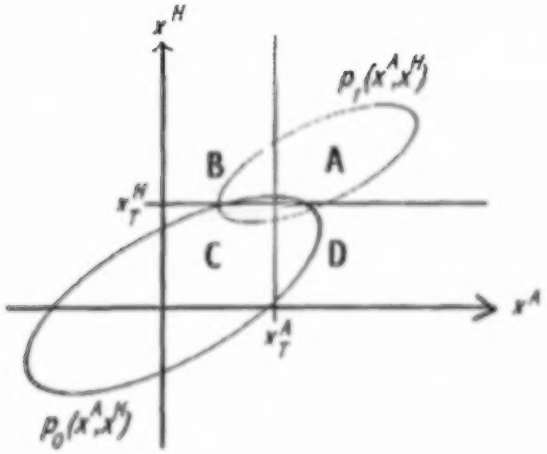

Figure 3: The combined ATR (A) and human (H) detection can be modeled using bivariate random distributions, one for the target class (1), and another for non-targets (0). Here the ATR decision variable is plotted on the horizontal axis, and the human operator's (effective) decision variable is plotted on the vertical. Each agent registers an alarm when their decision variable exceeds their threshold. The two thresholds divide the plane into four quadrants, labeled A, B, C, and D here.

Quadrant	Operator alarm	ATR alarm
A	yes	yes
B	yes	no
C	no	no
D	no	yes

Table 1: The straight-line decision boundaries divide the decision plane (Fig (3)) into four quadrants. Alarms are raised by each agent for an event (target or non-target) depending on the quadrant in which the event falls.

the operator to take x^A into account, however, through a simplified low-medium-high "confidence" ranking perhaps, in this way making the best under the circumstances. But such a system would also require an unprecedented degree of operator training, for the operator must now learn to systematically and optimally weigh the alarm threat level as one factor among many visual cues when making an alarm decision. Failing that, the decision boundaries are those of the operator and ATR considered independently, and appear as vertical lines in Figure (3). That is the assumption here, although the analysis could be extended to sloping linear, or non-linear decision boundaries if warranted.

The straight-line decision boundaries divide the (x^A, x^H) plane of Figure (3), and hence each class of objects, into four sub-classes in accordance with the alarms they raise, as summarized in Table 1.

The four fusion rules listed earlier prescribe how these quadrants are to be combined to

P517270 PDF [Page: 26 of 40]

raise a single alarm from the combined human-ATR system.

5.1 Joint performance

The prospective performance when the two agents are combined, when all other benefits of penalties of the ATR and on-screen alarming are for simplicity held equal, is given by

Fusion Rule	Quadrant	Performance Probabilities
1 & 2	A & B	$\left.\begin{matrix} P_t^1 \\ P_t^2 \end{matrix}\right\} = \int_{x_T^H}^{\infty}\int_{-\infty}^{\infty} p_t\left(x^A, x^H\right) dx^A dx^H = P_t^H$
3	A, B, & D	$P_t^3 = P_t^H + \int_{x_T^H}^{\infty}\int_{x_T^A}^{\infty} p_t\left(x^A, x^H\right) dx^A dx^H$
4	A	$P_t^4 = \int_{x_T^H}^{\infty}\int_{x_T^A}^{\infty} p_t\left(x^A, x^H\right) dx^A dx^H$

(2)

Fusion rule 1 gives much the same performance as operator working alone because the operator overrules every ATR alarm. It is possible that the ATR alarms incline an operator to register targets that would have otherwise passed without note in the absence of the alarm. This is equivalent to a downward slide in the operator's threshold x_T^H from quadrant A down into quadrant D, with no corresponding slide from B down to C. Such a slide may or may not improve overall performance, depending on the disposition of the joint p.d.f.'s relative to the thresholds.

Rule 2 likewise gives much the same performance of as operator working alone, assuming now that the operator is capable of inspecting and ruling on every ATR alarm. If the operator becomes overloaded, then some uninspected events in quadrant D will be pass as alarms because the operator was unable to reject them. A downward sliding of x_T^H between quadrants A and D is also possible. This again may or may not improve overall performance, depending on the disposition of the joint p.d.f.'s.

Fusion rules 3 and 4 can influence the overall detection performance dramatically. If the thresholds $\left(x_T^A, x_T^H\right)$ during fusion are the same as for independent operation, then the OR rule increases both the probability of detection and false alarms of the independent agents

$$P_t^3 \geq \max\left(P_t^H, P_t^A\right); \tag{3}$$

whereas the AND rule decreases both

$$P_t^4 \leq \min\left(P_t^H, P_t^A\right), \tag{4}$$

But the thresholds may of course be changed for fusion. The condition (3) can sometimes be overcome, for instance, by adjusting the ATR to operate with a much

lower false alarm rate than it would if it were working on its own; and condition (4) by adjusting it to operate with a much higher probability of detection than during autonomous operation. Re-adjustment of the ATR is generally necessary to optimize the fused performance, as we shall see later in Figure (4) in the example below, section 6.2.

5.2 Operator Workload

Mines are rare in practice, so the operator is occupied almost entirely with clutter and false alarms. For workload calculations we therefore consider the non-target class (0) alone. Let us say that the operator's tasks include 1) visually inspecting signatures on the screen, to which we assign one "unit of work" (UOW); and 2) the registration of an alarm when deemed necessary, presumably by positioning a cursor over the target and pressing a button, to which we shall assign 2 UOW, somewhat more effort than inspection, the precise relative value not being critical for demonstration purposes. These two tasks are summarized in the first row Table 2. The workload W associated with a task is given by

$$W = w \int\int_{Q} p_0\left(x^A, x^H\right) dx^A dx^H, \tag{5}$$

in which w is the unit of work, and Q represents one or more of the quadrants in Figure (3), associated with the task. Thus, operating without automation, the operator's relative workload is

$$\begin{aligned} W &= 1 \times \int\int_{A,B,C,D} p_0\left(x^A, x^H\right) dx^A dx^H + \\ &\quad 2 \times \int\int_{A,B} p_0\left(x^A, x^H\right) dx^A dx^H \\ &= 1 + 2P_0^H \end{aligned} \tag{6}$$

We take this for our reference work level against which the CAD system must be compared. Note that it is not an absolute workload because no clutter event rate has been specified. None is required for making relative workload comparisons; we need only assume that the clutter event rate is the same throughout.

Workload reductions occur when the number of class-0 events to which the operator must respond is decreased; the reduction being proportional to the decrease in the area in quadrants A and B, whose events must be handled by the operator. Workload penalties due to alarm masking and the need for repeated ATR adjustments are too system-specific to be included here, but they must be kept in mind as potentially significant burdens for the operator.

We turn now to the CAD system, with ATR operating under the four different fusion rules summarized in Table 2.

For fusion rule 1, the operator tasks are 1) to visually inspect the entire image as before; 2) raise missed alarms in quadrant B, and 3) acknowledge ATR alarms in quadrant A, which we assume to be an effort equal to raising an alarm (2 UOW). This amounts to the same workload as without ATR. Hence this most common fusion rule does not

Mode	Quadrants	Task	UOW
No ATR	A,B,C,D	visual inspection	1
	A,B	raise alarm	2
Rule 1	A,B,C,D	visual inspection	1
	A	acknowledge alarm	2
	B	raise missed alarm	2
Rule 2	A,B,C,D	visual inspection	1
	D	reject alarm	2
	B	raise missed alarm	2
Rule 3	B,C	visual inspection	1
	B	raise missed alarm	2
Rule 4	A,D	visual inspection	1
	A	acknowledge alarm	2

Table 2: Operator workload estimates under each fusion rule. Operator tasks are apportioned according to the operator's and ATR's inclination to raise alarms, that is, by the quadrants (column 2) in which the events fall; the quadrants being defined in Fig.(3). "UOW" stands for units of work as defined in the text.

decrease operator workload, but actually increases it when the overhead costs of alarm masking and ATR adjustments are included. As we saw earlier, moreover, this rule also provides no independent-agent gains because the ATR alarms are entirely subject to the operator's judgement. Thus there remains just one avenue by which such a CAD system could improve performance under rule 1, and that is by improving operator vigilance, goal (*c*) above, which has nonetheless received little attention from CAD system design community. Such confusion of purpose in the design of CAD may account for much of the operator frustration with existing systems, why they are likely to find them burdensome and of questionable benefit.

For fusion rule 2, the operator's tasks are 1) to visually inspect the entire image as before; 2) raise missed alarms in quadrant B; and 3) to reject false ATR alarms in quadrant D, a task comparable to raising an alarm, to which we likewise ascribe 2 *UOW*. There is clearly a potential for workload reduction inasmuch as the ATR now raises all alarms in quadrant A which were formerly raised by the operator.

For fusion rule 3, "OR", assuming that the ATR alarms are switched on, the operator's tasks are 1) to visually inspect the unalarmed events in quadrants B and C (because those in quadrants A and D have already been alarmed by the ATR), and 2) to raise missed alarms in quadrant B. Notice that, if the overall system performance goal of keeping the probability of false alarm low is ignored, then it would be possible to reduce the operator's workload to zero by sliding the ATR detection threshold ever lower, making the contributions from quadrants B and C essentially zero. Thus a balance must be struck between workload reduction and overall performance under this

P517270.PDF [Page: 29 of 40]

rule. If the on-screen alarms are switched off, as they might be under rule 3, then there is of course no reduction in operator workload.

For fusion rule 4, "AND", again assuming that the ATR alarms are switched on, the operator's tasks are 1) to visually inspect only the alarmed events in quadrants A and D (because the operator cannot raise alarms to which the ATR does not agree under this rule), and 2) to acknowledge the alarms deemed to be true in quadrant A. However radical this approach may first appear, it nevertheless supports the most dramatic reduction in operator workload, and may prove more tolerant of poor quality ATR algorithms than the other rules are; the high false alarm rates typical of poor ATR being more than offset by the overall reduction in operator workload. One could in fact design ATR for simply registering regions of interest (ROI) alone, without pressing on to automated discrimination between objects.

Changes in operator workload owing the ATR are illustrated in the example at the end of this paper, with related figures (5) to (7).

5.3 Operator Vigilance

We offer the following tentative observations regarding vigilance from work still in progress.

Vigilance can of course be kept high by rotating operators through short watches, which is good practice for any critical detection task. The optimal rate of rotation presumably increases as the operator workload increases, whether due to higher clutter rates or the management of the CAD system. Beyond watch rotation, however, the ATR on-screen alarms can also increase vigilance, by providing a momentary diversion which raises the operator's arousal, and reminds them of the target-like signatures to watch for [3]. The ATR alarms must therefore be occasional, and at a rate that is not entirely determined by the amount of clutter. Designers might consider using different levels of alarms, such as a more common "low-level" alarm and a rare "high"-level alarm, to much the same effect. But the implications for system design can be even more radical.

If vigilance gains call for low probabilities of false alarm, then, in light of the detection sensitivities now achieved by existing ATR, it would appear that one must also accept low probabilities of detection; significantly lower, that is, than either the probability of detection expected of the human operator alone, of the human and ATR working together in the CAD system, or of the ATR working alone for that matter. It is the change in role, from autonomous operation to operator assistant, that calls for the readjustment of the ATR thresholds on integration in the CAD system, in this way maximizing the overall human-machine performance through vigilance gains more so than in any other way. This is especially true of the most widely used fusion rule 1, which looks finally to gains in operator vigilance for its justification.

Rule 2, on the other hand, calls for a balance between workload reduction and operator vigilance inasmuch as the ATR thresholds minimizing workload may differ from those

optimizing vigilance. Under this rule it is perhaps best to let the operator adjust the ATR sensitivity as they feel inclined, either to replicate their judgements as much as possible for greatest workload reduction, or to increase x_T^A for a lower false alarm rate for greatest vigilance increase.

Rule 3 ("OR") offers potential independent-agent gains and workload reduction, naturally requires lower ATR false alarm rates (3), making it a good candidate for vigilance gains.

Rule 4 ("AND") also offers potential independent-agent gains and workload reduction, but naturally requires high ATR false alarm rates (4). Vigilance gains might nevertheless be an indirect result of the potentially dramatic workload reductions noted above for this rule.

6. Example

For brevity we assume that the ATR performs roughly as well as the human operator—i.e., they have roughly the same ROC curves. The operator's level of performance is a goal to which ATR aspires that has not yet been achieved. The equal-performance assumption is therefore optimistic, and projects forward into the future, in the direction that ATR is presumably headed. This is not to say that the operator and ATR make the very same decisions. It is a statistical equivalence, not a decision-by-decision equivalence.

The performance of each fusion rule can be demonstrated in more detail when the joint p.d.f.'s $p_{1,0}\left(x^A, x^H\right)$ are known. Ultimately these must be determined experimentally, which is non-trivial owing to the costly requirements for many randomized trials to eliminate bias. And the results are of narrow application: to the level of operator skill and the particular ATR algorithms used in the trials. In a preliminary study it is better to assume a plausible probability density distribution, and proceed from there.

6.1 Bivariate Gaussian distribution

As in the SDT, we therefore assume the bivariate normal distribution [1],

$$p_i\left(x^A, x^H\right) = \frac{1}{\sqrt{2\pi\left|\Sigma_i\right|}} \exp\left[-\frac{\left(X - M_i\right)^T \Sigma_i \left(X - M_i\right)}{2}\right] \tag{7}$$

in which $i = 1, 0$ for the target and non-target classes, respectively; $X = \left[x^A, x^H\right]^T$ is the decision vector; $M_i = \left[\mu_i^A, \mu_i^H\right]^T$ is a vector holding the means for the i'th class; and Σ_i is a symmetric 2X2 covariance matrix,

$$\Sigma_i = \begin{bmatrix} \nu_i^A & \nu_i^{AH} \\ \nu_i^{AH} & \nu_i^H \end{bmatrix}; \tag{8}$$

P517270.PDF [Page: 31 of 40]

ν^A and ν^H being the independent agent's variances, and ν^{AH} being the correlation of the decision variables. Justification for assuming the normal distribution for the human operator is given in the references [1] [3]. Justification for ATR follows straightforwardly from the central limit theorem if its algorithms are based on linear discriminators [10], and more generally, insofar as random noise, target orientation, glints, and local seafloor properties all affect target recognition simultaneously, each with independent randomness. The *sensitivity* of each agent considered separately can be defined as their relative class separation in the decision variable space

$$d^{A,H} = \frac{\mu_1^{A,H} - \mu_0^{A,H}}{\sqrt{\frac{1}{2}\left(\nu_1^{A,H} + \nu_0^{A,H}\right)}}. \tag{9}$$

The larger the sensitivity, the better the detection agent.

To simulate an ATR designed to discriminate on the basis of visual cues, such as highlight-shadow matching, object size estimates and so forth, much as an operator presumably does, we assume that there is a significant correlation between the ATR and operator, more so for the target class than the clutter class. In particular, for the bivariate Gaussian (7) we assume

$$M_1 = \begin{bmatrix} 1.5 \\ 1.5 \end{bmatrix} \text{ and } \Sigma_0 = \begin{bmatrix} 1.0 & 0.8 \\ 0.8 & 1.0 \end{bmatrix}. \tag{10}$$

$$M_0 = \begin{bmatrix} 0.0 \\ 0.0 \end{bmatrix} \text{ and } \Sigma_0 = \begin{bmatrix} 1.3 & 0.5 \\ 0.5 & 1.3 \end{bmatrix}. \tag{11}$$

In this case the detector sensitivities considered separately are

$$d^A = d^H = 1.293, \tag{12}$$

which is realistic insofar as $P_1^H \approx 0.80$ for detection, for example, and $P_0^H \approx 0.3$ is realistic for a human operator, discriminating between mine-sized clutter objects and mines say, to pick just one possible operating point from the ROC curve (see Fig.(4)). The operator may of course be more or less inclined to raise an alarm (depending on the impression they have of the likelihood of encountering a target, and of the costs of making an error), which in turn determines their threshold x_T^H. The ATR works at some other point along the ROC curve, according to its threshold x_T^A, presumably determined by the CAD system designer so as to optimize the assisting role of the ATR.

6.2 CAD system performance

In equation (5.1) we saw, from a purely signal processing point of view, ignoring workload and vigilance factors, that no performance changes are expected under fusion rules 1 and 2, because the ATR alarms are subordinated to the operator. Thus the ROC

P517270.PDF [Page: 32 of 40]

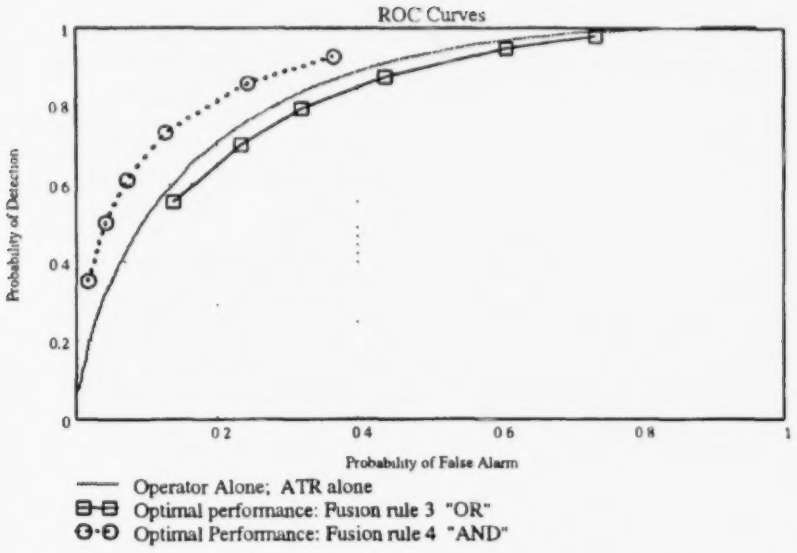

Figure 4: ROC curves for equal operator and ATR detection assuming the bivariate normal distributions given in (11) and (12). The plain solid line is the ROC curve for the operator and ATR operating independently, assumed to be equal here for simplicity. The squares mark the best possible performance under fusion rule 3, the "OR" rule. The circles mark the best possible performance under rule 4, the "AND" rule. In this case rule 4 clearly outperforms rule 3. But this is not generally true. The reverse may be true, depending on the disposition of the class distributions (Fig (3)). Note that the optimal performance under either rule is only achieved when the operator and the ATR adjust their thresholds in concert. The optimization was achieved here by numerical iteration. It is not clear how the optimal would be achieved in practice, so these curves represent the ultimate performance limits under ideal conditions, for the assumed density functions.

curves for the overall CAD system are roughly those of the operator working independently. We therefore turn to rules 3 and 4.

Figure (4) shows the best possible performance, when one or both of the operator and ATR adjust their thresholds to give the optimal fusion, found here by maximizing detection $P_1^{3,4}$ in $\left(x_T^A, x_T^H\right)$ for a given fixed false alarm probability $P_0^{3,4}$. Optimization is much more difficult in practice, if not impossible, when the joint p.d.f. is not known, but here we have a rough estimate of the best performance improvements one could typically expect to gain from such optimal fusion, owing to the joint cooperation of independent detection agents [goal (a) above], if such optimization were possible, from which we can judge whether such gains justify the efforts required to optimize in this respect.

P517270.PDF [Page: 33 of 40]

6.3 Operator Workload

In section (5.2), we saw that fusion rule 1 offered no workload reduction. The workload is that of operating without ATR, plus the burden of alarm masking and ATR adjustments.

Figure (5) shows the change in operator workload (defined in section 5.2) induced by fusion rule 2. The vertical axis shows the change in workload (UOW), relative to operation without ATR, expressed in percent,

$$\frac{\text{workload with ATR} - \text{workload without ATR}}{\text{workload without ATR}} \times 100\ \%. \qquad (13)$$

The horizontal axis is the detection threshold for the ATR, which can in principle be adjusted to any value, though the implications for overall performance will have a bearing on the range of acceptable values. Thus low threshold values result in high ATR false alarm rates, for instance, whereas very high threshold values (greater than 2.5 say) effectively shut the ATR off altogether, allowing very few alarms. Of particular note in the figure are the minima in workload, which occur, as we might expect, near points where the ATR threshold equals that of the operator. This is indeed what the operator's adjustment to the ATR (section (2.3.1)) presumably achieves, at least insofar as decision-by-decision matching is possible between the operator and ATR. For fusion rule 2, then, we have potentially significant reductions in workload, depending on the disposition of the joint p.d.f.'s (Fig. (3)).

Figure (6) shows the change in workload for fusion rule 3, "OR". Workload decreases as the ATR threshold decreases because the ATR misses fewer of the alarms that the operator would raise. The overall system performance goals, of minimizing the probability of false alarm oppose the inclination to decrease the ATR threshold very much.

Figure (7) shows the relative change in workload for fusion rule 4, "AND". The reduction in workload is potentially dramatic, though exaggerated somewhat in the figure. The elimination of work to the right occurs because the ATR is effectively shut off by the high threshold, and is no longer registering any alarms, which would in fact result in near zero probability of detection for the overall CAD system. Thus, here, as with all Figures (5) to (7), the overall performance constraints must be kept in mind as additional constraints. A more detailed analysis is required to investigate the performance-workload trade-offs in for fusion rules.

7. Conclusions

The two main points made here apply to CAD systems quite generally, and are perhaps a reversal in the perspective of CAD design for those who have so far avoided questions about human factors. The first is that the detection performance of the automatic target recognition (ATR), when serving as a sub-system designed to maximize the overall

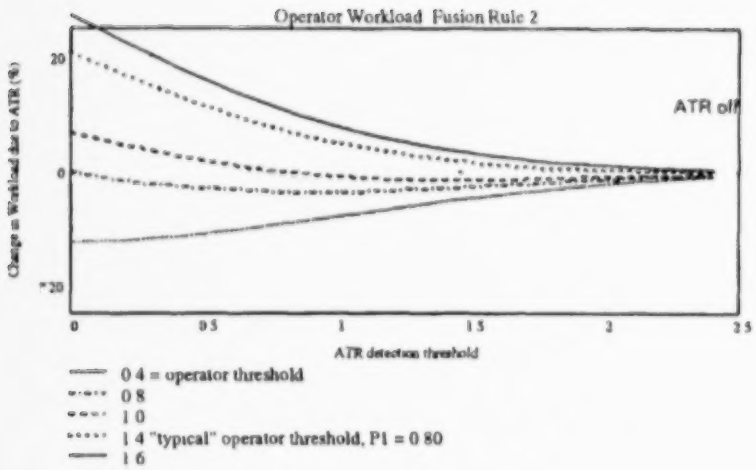

Figure 5: The relative change in operator workload under fusion rule 2, as a function of the ATR detection threshold. Negative changes signify reductions in workload. In this case, the anticipated change in operator workload would not be very significant, less than 5 % for realistic ATR thresholds greater than 1.0.

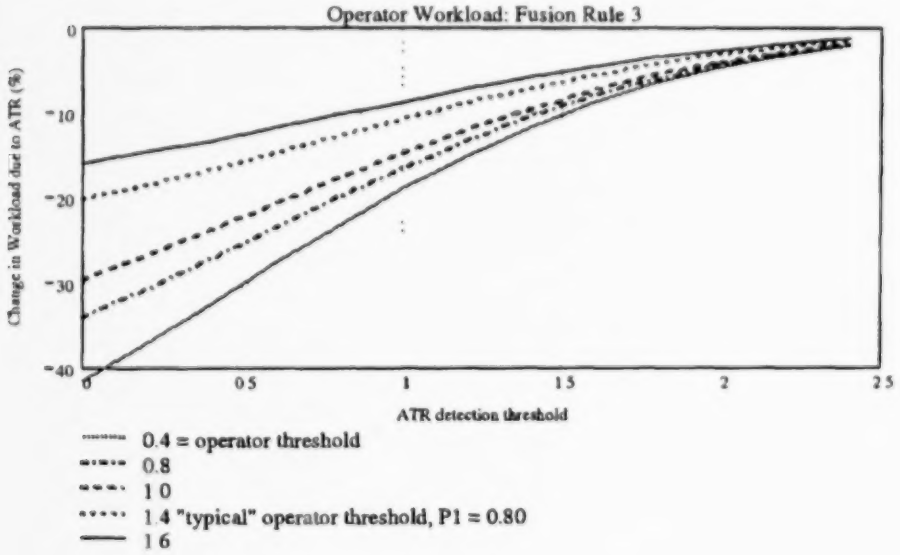

Figure 6: The relative change in operator workload under fusion rule 3, as a function of the ATR detection threshold. Unlike rule 2, this rule always reduces operator workload. In this example, the reduction is more significant as well, perhaps a 10 % reduction for ATR thresholds greater than 1.0.

P517270.PDF [Page: 35 of 40]

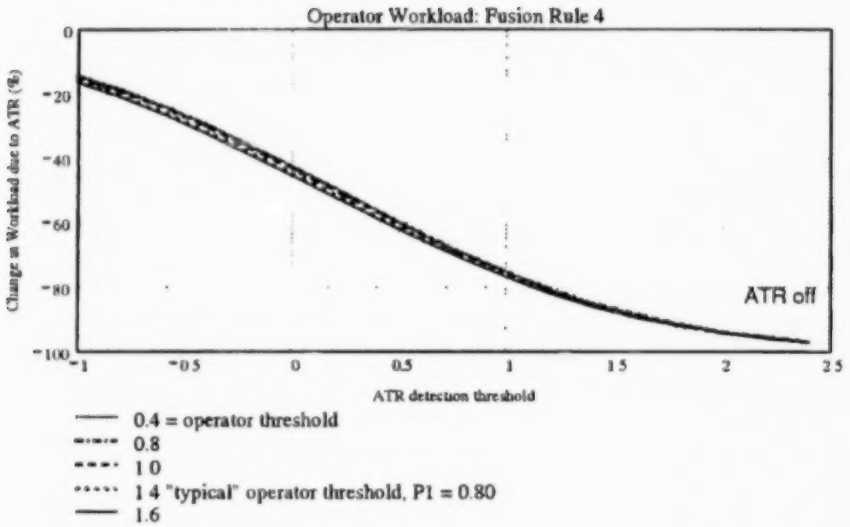

Figure 7: *The relative change in operator workload under fusion rule 4, as a function of the ATR detection threshold. This rule affords the greatest potential for workload reduction, perhaps a 75 % reduction for ATR thresholds near 1.0. A more detailed analysis would be required to determine if the performance at this workload reduction would be acceptable. Even a 20 % reduction near ATR thresholds of -1.0 would be greater than those encountered with the other fusion rules in this example*

performance of the human-computer team, is likely to be significantly different from the performance achieved by the team itself. This is because the objectives of the ATR in a CAD system are no longer to *directly* minimize the risks and costs of mission operation (as it would be in autonomous operation), but to *indirectly* minimize them by serving as an aid to the operator—a task which entails quite another set of risks and costs, many of which are purely human performance factors, workload and vigilance, as described here. Thus, although the ATR sub-system may be designed and tested as if it were intended for autonomous operation, with a view to maximizing its *sensitivity* (as defined above), its operation as a subordinate member of the human-computer team calls for readjustment of its thresholds to maximize its contribution as an aid.

The second point is that clear system performance goals, and a clearly defined mode of operation (fusion rules), are absolutely crucial for the purposeful design of CAD systems. A cost-benefit or performance analysis of any kind would in fact be impossible without them. Much of the dissatisfaction and confusion presently surrounding CAD systems no doubt comes from the lack of clarity in either respect. Three overall performance goals and four fusion rules were identified here.

Of the fusion rules, the first and most cautious rule—cautious inasmuch as it prohibits all autonomous ATR decisions—is the most widely used at present. The remaining three fusion rules were offered as possible alternatives because they grant the ATR some degree of autonomy for raising automated alarms, although these rules have apparently been avoided in practice for precisely that reason. Greater autonomy might be also be achieved under other fusion rules that were not considered here, such as weighted voting between the ATR and human operator, and an analysis similar to that given here could be applied to assess those rules as well. Greater ATR autonomy seems to be the only way to realize significant gains in performance that capitalize on its ability to "see" more thoroughly and with higher resolution and precision than an operator, and to simultaneously reduce the operator's workload. Future CAD system developments should therefore include plans in that direction.

References

1. Green, D.M., and Swets, J.A., "Signal detection theory and psychophysics," Peninsula Publishing, Los Altos, California, 1988.

2. Wickens, C.D., Gordon, S.E., and Liu, Y., "An introduction to human factors engineering," Addison Wesley Longman Inc., New York, 1998.

3. Wickens, C.D, and Hollands, J.G., "Engineering psychology and human performance," 3rd Edition, Prentice Hall, New Jersey, 2001.

4. Kessel R.T., "Mission-adaptive computer aided detection (CAD) for mine hunting: system concept and theory," Canadian National Defence Research and Development Branch Technical Memorandum DREA TM 2000-111, Feb 2001.

5. Wort, P., CAD/CAC Study Report, MacDonald Dettwiler Report CL-RP-50-7394, Issue 1, Revision 0, June 27 1996.

6. Pratt, W.K., "Digital image processing," 2nd Edition, John Wiley and Sons, New York, 1991.

7. Van Trees, H.L., *Detection, estimation, and modulation theory*, John Wiley and Sons, New York, 1968.

8. Fukunaga K., "Introduction to statistical pattern recognition", 2nd Ed., Academic Press, New York, 1990.

9. Devroye L. , Gyorifi L. , and Lugosi G., "A probabilistic approach to pattern recognition", Springer, New York, 1996.

10. Kessel R.T., "A probabalistic theory for the design of optimal linear discriminators for the automated detection of objects in sidescan sonar images," Canadian National Defence Research and Development Branch Technical Memorandum DREA TM 2000-021, Feb. 2000.

P517270.PDF [Page: 38 of 40]

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1 ORIGINATOR (the name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8.)

Defence Research Establishment Atlantic
PO Box 1012
Dartmouth, Nova Scotia, Canada B2Y 3Z7

2 SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable)

UNCLASSIFIED

3 TITLE (the complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title)

On-Screen Alarms in Computer-Aided Detection Systems: *Combining Signal Detection, Human Factors, and System Design*

4 AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)

Ronald T. Kessel

5 DATE OF PUBLICATION (month and year of publication of document)	6a NO. OF PAGES (total containing information. Include Annexes, Appendices, etc.)	6b NO. OF REFS (total cited in document)
Nov 2001	34 (approx.)	10

7 DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered)

Technical Memorandum (TM)

8 SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development. Include address)

Defence Research Establishment Atlantic

9a PROJECT OR GRANT NO. (if appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant)	9b CONTRACT NO. (if appropriate, the applicable number under which the document was written)
1da12	

10a ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b OTHER DOCUMENT NOs. (Any other numbers which may be assigned this document either by the originator or by the sponsor.)
DREA TM 2001-184	

11 DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
( ) Defence departments and defence contractors; further distribution only as approved
( ) Defence departments and Canadian defence contractors; further distribution only as approved
( ) Government departments and agencies; further distribution only as approved
( ) Defence departments; further distribution only as approved
( ) Other (please specify):

12 DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11)) is possible, a wider announcement audience may be selected)

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM

P517270.PDF [Page: 39 of 40]

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual)

It is usual for computer-aided detection (CAD) systems to report automated alarms by highlighting their position on the operator's primary viewing screen. Here we carry out a "cost-benefit" analysis of the practice, contrasting its operational utility and liabilities. Because on-screen alarms are so very central to CAD systems, a cost-benefit analysis provides a framework for a critique of CAD systems generally, combining aspects of signal detection theory, human factors, and system design. The analysis highlights the shortcomings of existing CAD systems, and identifies the source of widespread dissatisfaction with them; why they are seen more often as a burden than a help by operators. The primary cause, it is argued, is that CAD systems are designed without a clear specification of the assisting role that automatic target recognition (ATR) is to fulfil for the operator. Clarified purpose in turn requires well-defined fusion rules by which the automated alarms are combined with the operator's, and by which the assisting goals are ultimately achieved. Given clarified purpose and fusion rules, it can be shown that the optimal integration of the ATR as an operator aid is likely to require radically different algorithm settings than if---as ATR developers generally assume for their algorithms---ATR were to operate autonomously. This may be a significant change in perspective for system designers who have so far avoided questions of human factors and overall performance. Non-conventional modes of CAD system operation are considered, showing how they offer different opportunities for assisting the operator.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title)

computer-aided detection, human monitored alarm systems, automatic target recognition, minehunting, signal detection theory